

Accepted to ApJ

Voigt-Profile Analysis of the Lyman-alpha Forest in a Cold Dark Matter Universe

Romeel Davé and Lars Hernquist¹

Astronomy Department, University of California, Santa Cruz, CA 95064

David H. Weinberg

Astronomy Department, Ohio State University, Columbus, OH 43210

and

Neal Katz

Astronomy Department, University of Washington, Seattle, WA 98195

ABSTRACT

We use an automated Voigt-profile fitting procedure to extract statistical properties of the Ly α forest in a numerical simulation of an $\Omega = 1$, cold dark matter (CDM) universe. Our analysis method is similar to that used in most observational studies of the forest, and we compare the simulations to recently published results derived from Keck HIRES spectra. With the Voigt-profile decomposition analysis, the simulation reproduces the large number of weak lines ($N_{\text{HI}} \lesssim 10^{13} \text{ cm}^{-2}$) found in the HIRES spectra. The column density distribution evolves significantly between $z = 3$ and $z = 2$, with the number of lines at fixed column density dropping by a factor ~ 1.6 in the range where line blending is not severe. At $z = 3$, the b -parameter distribution has a median of 35 km s^{-1} and a dispersion of 20 km s^{-1} , in reasonable agreement with the observed values. The comparison between our new analysis and recent data strengthens earlier claims that the Ly α forest arises naturally in hierarchical structure formation as photoionized gas falls into dark matter potential wells. However, there are two statistically significant discrepancies between the simulated forest and the HIRES results: the model produces too many lines at $z = 3$ by a factor $\sim 1.5 - 2$, and it produces more narrow lines ($b < 20 \text{ km s}^{-1}$) than are seen in the data. The first result is sensitive to our adopted normalization of the mean Ly α optical depth, and the second is sensitive to our assumption that helium reionization has not significantly raised gas temperatures at $z = 3$. It is therefore too early to say whether these discrepancies indicate a fundamental problem with the high-redshift structure of the $\Omega = 1$ CDM model or reflect errors of detail in our modeling of the gas distribution or the observational procedures.

¹Presidential Faculty Fellow

Subject headings: galaxies: formation — large-scale structure of universe — line: profiles — methods: numerical — quasars: absorption lines

1. Introduction

Absorption lines in quasar spectra, especially the “forest” of Ly α lines produced by concentrations of neutral hydrogen, are uniquely suited for probing structure formation in the high-redshift universe. The absorbers trace relatively pristine baryonic material over a wide range of redshifts, densities, temperatures, and ionization states. Recent advances in computer technology have expanded our ability to predict the conditions of the absorbing gas, while high-precision observations made using the HIRES spectrograph (Vogt, et al. 1994) on the 10m Keck telescope have quantified the statistics of the low column density absorbers to unprecedented accuracy (e.g., Hu et al. 1995, hereafter HKCSR). These data provide stringent constraints on theories of structure formation and the state of the intergalactic medium (IGM) at high redshifts. However, a detailed confrontation between theory and observations requires that simulations and observed quasar spectra be analyzed using similar techniques. In this paper we apply an automated Voigt-profile fitting algorithm to Ly α spectra from a simulation of the cold dark matter (CDM) scenario (Peebles 1982; Blumenthal et al. 1984). We compare the statistics of the resulting line population to those derived by HKCSR from HIRES spectra.

Recent cosmological simulations that incorporate gas dynamics, radiative cooling, and photoionization reproduce many of the observed features of quasar absorption spectra, suggesting that the Ly α forest arises as a natural consequence of hierarchical structure formation in a universe with a photoionizing UV background (Cen et al. 1994; Zhang, Anninos, & Norman 1995; Hernquist et al. 1996, hereafter HKWM; Katz et al. 1996b; Miralda-Escudé et al. 1996). Most of the low column density lines are produced by structures of moderate overdensity that are far from dynamical or thermal equilibrium, blurring the traditional distinction between the Ly α forest and Gunn-Peterson (1965) absorption from a smooth IGM. HKWM used a simple flux threshold algorithm to identify lines in their simulated spectra, defining any region with transmitted flux continuously below a specified threshold as a single line. They showed that their simulation of a CDM universe reproduced the observed abundance of absorption systems as determined by Petitjean et al. (1993, hereafter PWRCL) quite well over most of the column density range $10^{14}\text{cm}^{-2} < N_{\text{HI}} < 10^{22}\text{cm}^{-2}$, with a significant discrepancy for $N_{\text{HI}} \sim 10^{17}\text{cm}^{-2}$.

The traditional technique for identifying and characterizing quasar absorption lines is to fit spectra by a superposition of Voigt profiles. The HIRES spectra have very high signal-to-noise ratio and resolution, and most of the lines found in this way are weak absorbers with column densities $N_{\text{HI}} < 10^{14}\text{cm}^{-2}$. The threshold and Voigt-profile procedures behave very differently in this regime, since a single feature identified by the threshold method will often be decomposed into

a blend of weaker lines when it is modeled as a superposition of Voigt profiles. In order to compare to published line population statistics from HIRES data, therefore, it is essential to analyze the simulated spectra by Voigt-profile decomposition.

The physical model implicit in the decomposition technique is that of a collection of discrete, compact clouds, each characterized by a single temperature (or at least by a single velocity dispersion, which could include contributions from thermal motion and from Gaussian-distributed “turbulent” velocities). The simulations undermine this physical picture because the absorbing systems merge continuously into a smoothly fluctuating background, often contain gas at a range of temperatures, and are usually broadened in frequency space by coherent velocity flows that do not resemble Gaussian turbulence. Nonetheless, any spectrum can be described phenomenologically by a superposition of Voigt-profile lines, with the number of components increasing as the signal-to-noise ratio improves and more subtle features must be matched. The distributions of fitted column densities and b -parameters provide a useful statistical basis for comparing simulations and observations, and this is the approach that we adopt in this paper. We will discuss the correspondence between the parameters of the Voigt-profile components and the physical state of the absorbing gas elsewhere (Davé et al., in preparation).

2. Simulation and Artificial Spectra

The simulation analyzed here is the same as that of HKWM: a CDM universe with $\Omega = 1$, $H_0 = 50 \text{ km s}^{-1}\text{Mpc}^{-1}$, baryon fraction $\Omega_b = 0.05$, and a periodic simulation cube 22.222 comoving Mpc on a side containing 64^3 gas particles and 64^3 dark matter particles, with individual particle masses of $1.45 \times 10^8 M_\odot$ and $2.8 \times 10^9 M_\odot$, respectively. The power spectrum is normalized to $\sigma_8 = 0.7$, roughly the value required to reproduce observed galaxy cluster masses (Bahcall & Cen 1992; White, Efstathiou, & Frenk 1993); we will consider COBE-normalized CDM models with $\Omega = 1$ and $\Omega < 1$ in future work. We use the N-body + smoothed-particle hydrodynamics code TreeSPH (Hernquist & Katz 1989) adapted for cosmological simulations (Katz, Weinberg, & Hernquist 1996, hereafter KWH) to evolve the model from $z = 49$ to $z = 2$.

Instead of the ν^{-1} UV background spectrum adopted by HKWM, we use the spectrum of Haardt & Madau (1996; hereafter HM), which is computed as a function of redshift based on the UV output of observed quasars and reprocessing by the observed Ly α forest. The spectral shape is significantly different from ν^{-1} , but the UV background influences our Ly α forest results primarily through the HI photoionization rate Γ , a cross-section weighted integral of the spectrum (KWH, equation 29). The new simulation also includes star formation and feedback (see KWH), but this has no noticeable effect on the Ly α forest results. A comparison between the galaxy populations of this simulation and the HKWM simulation appears in Weinberg, Hernquist, & Katz (1996).

The mean opacity of the Ly α forest depends on the parameter combination Ω_b^2/Γ . Since observational determinations of Ω_b and Γ remain quite uncertain, we treat the overall intensity

of the UV background as a free parameter and scale it to match the observed mean Ly α optical depth $\bar{\tau}_\alpha$. When evolving the simulation, we divide HM’s intensities by a factor of two, retaining their redshift history and spectral shape. We find that we must reduce the intensities by further factors of 1.28 and 1.38 at $z = 2$ and $z = 3$, respectively, in order to match the estimate $\bar{\tau}_\alpha = 0.0037(1+z)^{3.46}$ of Press, Rybicki, & Schneider (1993; hereafter PRS). Although we apply this final reduction only at the analysis stage, to compute neutral fractions when generating spectra, the result is virtually identical to that of changing the intensity during dynamical evolution (Miralda-Escudé et al. 1997, hereafter MWHK). In order to match the PRS mean optical depth at $z = 3$ with the original HM background intensity we would need $\Omega_b \approx 0.08$, closer to the value advocated by Tytler, Fan, & Burles (1996; but see Rutgers & Hogan 1996 and references therein). The value of $\bar{\tau}_\alpha$ plays the role of a normalizing constraint, used to fix the important combination of free parameters in our IGM model. Once Ω_b^2/Γ is set, there is no further freedom to adjust the simulation predictions, and the remaining properties of the Ly α forest provide tests of the cosmological scenario itself.

We generate artificial spectra at $z = 2$ and $z = 3$ along 300 random lines of sight through the simulation cube, using the methods described in HKWM and Cen et al. (1994). We do not consider higher redshifts here because the strong absorption leads to severe blending of lines. The wavelength spread across the box is 23.4 Å at $z = 2$ and 35.9 Å at $z=3$. Each artificial spectrum contains 1000 pixels; an individual pixel has a velocity width $\sim 2 \text{ km s}^{-1}$ and a spatial extent ~ 20 comoving kpc, twice the gravitational softening length. In the Ly α forest regime, the gas distribution is smooth on these scales.

3. Fitting Voigt Profiles to Artificial Spectra

We want the analysis of our simulated spectra to closely match that used in typical observational studies, HKCSR in particular. To this end, we have developed an automated Voigt-profile fitting routine, AUTOVP, which allows us to efficiently handle large quantities of simulated data and which provides an objective algorithm that can be applied to observational data.

We add noise to our simulated spectra employing a combination of Gaussian photon noise with signal-to-noise ratio $S/N = 50$ in the continuum (corresponding roughly to the HKCSR data) and a fixed readout noise chosen to match the characteristics of the Keck HIRES spectrograph. Varying S/N changes our results only at the lowest column densities. While we know the true continuum level in the simulated spectra, this is not the case for the observational data. We therefore estimate the continuum in the simulated spectra by the iterative procedure commonly used for Echelle data: fitting a third-order polynomial to the data set, excluding any points lying $\gtrsim 2\sigma$ below the fit, refitting the non-excluded points, and repeating until convergence is achieved. This technique is based upon the implicit assumption that the regions of lowest absorption in a high-resolution spectrum lie close to the true continuum level. Because the simulated spectra show

fluctuating Gunn-Peterson absorption that increases with z , continuum fitting has a systematic tendency to remove flux, an average of 1.2% at $z = 2$ and 5.7% at $z = 3$. The effect would probably be somewhat smaller in observational data because a typical HIRES Echelle order ($\sim 45\text{\AA}$) is longer than one of our simulated spectra, giving a higher probability that the spectrum contains a region of genuinely low absorption.

Given a normalized, continuum-fitted spectrum and its noise vector, we apply AUTOVP to detect lines and fit Voigt profiles. In its first phase, AUTOVP identifies lines and makes an initial estimate of their column densities and b -parameters. Detection regions are identified above an 8σ confidence level, following the method of Lanzetta, Turnshek, & Wolfe (1987). For line identification purposes, the data is convolved with a two-pixel-width Gaussian and 1σ of noise is subtracted to yield a “minimum flux”. For non-saturated regions, a single Voigt profile is placed at the lowest flux value in the detection region, and N_{HI} and b are reduced by small increments from large initial values until the model is everywhere above the minimum flux. For saturated regions a line is placed in the middle of the trough, and N_{HI} and b are adjusted to fit the “cusp” regions, *i.e.* regions about five pixels wide on either side of the trough. The resulting first-guess line is then subtracted from the data to obtain a residual flux. Detections regions are identified in the residual flux, and the procedure is repeated until there are no more 8σ detections. The line identification procedure is very robust, never failing for non-saturated regions and only occasionally producing a bad fit even in complex, blended, saturated regions, where N_{HI} and b are largely degenerate and the cusp regions are difficult to identify.

In its second phase, AUTOVP takes the initial guess and performs a simultaneous χ^2 -minimization on the parameters ($v_{\text{central}}, N_{\text{HI}}, b$) of all lines within each detection region. Three independent minimization techniques are employed in conjunction in order to reliably identify the global χ^2 minimum. AUTOVP then tries to remove any components with a formal error in N_{HI} or b comparable to the parameter value, refitting the detection region with one less line. If the resulting χ^2 is lower than the original value the rejection is accepted, otherwise the fit returns to the original set of lines. AUTOVP thus attempts to fit the spectrum with as few lines as possible while still minimizing χ^2 . If the fit after these line rejections is “good” (characterized empirically by $\chi^2 \lesssim 2$ per pixel), the program ends, otherwise it tries to add a line where the local contribution to χ^2 is greatest. The rare cases where AUTOVP fails to find a good fit are flagged for possible manual intervention. AUTOVP is designed to interface with the PROFIT interactive Voigt-profile fitting package (Churchill 1996). This graphical interactive fitter can be used to manually adjust poor fits, although this was required in so few cases that we base our simulation statistics entirely on the automated fits.

In Figure 1 we show the results of AUTOVP applied to a $S/N = 50$, continuum-fitted spectrum at $z = 3$. The bottom panel shows the first-guess fit superimposed on the simulated spectrum. The top panel shows the final fit after the χ^2 -minimization has been performed. Generally AUTOVP has greatest difficulty in obtaining first-guess fits in blended saturated regions like the one illustrated here. Nevertheless, the minimization produced an adequate fit with no

interactive adjustment. Lines with $N_{\text{HI}} \geq 10^{13}$ are indicated by the long tick marks above the spectrum, while lines with $N_{\text{HI}} < 10^{13}$ have short tick marks. The number of these small lines identified by AUTOVP is somewhat sensitive to the adopted S/N and the detection threshold. To keep our results fairly robust against details of our fitting procedure, we exclude these lines from our analysis and only focus on lines with $N_{\text{HI}} \geq 10^{13}$. AUTOVP has also been applied to observational data from H1216 and MgII2796, with results quite similar to those obtained from manual fitting.

4. Results

Figure 2 shows the column density distribution $f(N_{\text{HI}})$, the number of lines per unit redshift per linear interval of N_{HI} . Solid and dashed lines show the simulation results from AUTOVP at $z = 3$ and $z = 2$, respectively. The dotted line shows $f(N_{\text{HI}})$ obtained using the HKWM threshold algorithm at $z = 3$, with a flux threshold of 0.7. As expected, the two methods yield similar results at high column densities, $N_{\text{HI}} \gtrsim 10^{14.5} \text{ cm}^{-2}$, but at lower N_{HI} AUTOVP deblends much more and finds many more lines. We find a similar trend at $z = 2$, though because of the reduced line crowding at lower redshift the agreement between the two methods extends down to $N_{\text{HI}} \sim 10^{14} \text{ cm}^{-2}$.

Filled and open circles in Figure 2 show the observational results of PWRCL and HKCSR, respectively. The two determinations of $f(N_{\text{HI}})$ agree well in their regime of overlap, $10^{13.6} \text{ cm}^{-2} \lesssim N_{\text{HI}} \lesssim 10^{14.3} \text{ cm}^{-2}$. The high S/N and resolution of the HIRES data allow HKCSR to detect much weaker absorption features, and their $f(N_{\text{HI}})$ continues to rise down to the lowest bin, $N_{\text{HI}} \sim 10^{12.5} \text{ cm}^{-2}$. Clearly a comparison between the simulations and HKCSR’s published line statistics must be based on Voigt-profile fitting, since the vast majority of their lines lie in the region where line blending causes large differences between this method and the threshold algorithm. We compute the HKCSR $f(N_{\text{HI}})$ directly from their published line list, with no corrections for “incompleteness.” HKCSR estimate such corrections from artificial spectra *assuming* an underlying model of randomly distributed, Voigt-profile lines with a power-law $f(N_{\text{HI}})$ and a specified distribution of b -parameters, and they conclude that their results are consistent with $f(N_{\text{HI}}) \propto N_{\text{HI}}^{-1.46}$ down to $N_{\text{HI}} \approx 10^{12.5} \text{ cm}^{-2}$, where the correction for incompleteness is a factor of four. If we applied the same correction factors to the simulation results, the derived column density distributions would also rise in a nearly power-law fashion instead of turning over at low column densities. But the simulations provide no a priori reason to expect Voigt-profile lines or a power-law $f(N_{\text{HI}})$, so we prefer to compare them directly to the data without trying to correct either for lines “lost” to blending.

The mean redshift of the HKCSR lines is $\bar{z} = 2.9$, so the closest comparison is to the $z = 3$ simulation results. To make this comparison more exact, we convolved the $z = 3$ artificial spectra to a resolution $\Delta\lambda = 0.06 \text{ \AA}$ ($\Delta v = 3.7 \text{ km s}^{-1}$) before analysis, which has the minor effect of removing some lines with $N_{\text{HI}} \lesssim 10^{13}$. When analyzed by Voigt-profile decomposition,

the simulation reproduces the large number of weak lines found in the HIRES spectra. In fact, the simulation overproduces the number of lines by a factor of $1.5 - 2$ in the column density range $10^{13} \text{ cm}^{-2} \lesssim N_{\text{HI}} \lesssim 10^{14} \text{ cm}^{-2}$, a discrepancy that we will return to in §5. The rolloff at low column densities is also somewhat different, but the results for the weakest lines are the most sensitive to the details of the fitting procedure and to the modeling of noise and spectral resolution, so we regard this difference as less significant.

In the regime where line blending is unimportant, $f(N_{\text{HI}})$ of the simulations drops by a factor of $\sim 1.5 - 2$ between $z = 3$ and $z = 2$. At low column densities $f(N_{\text{HI}})$ actually increases because of the reduced effects of line blending. As discussed in HKWM, Miralda-Escudé et al. (1996), and MWHK, the evolution of the line population over this redshift range is driven primarily by the expansion of the universe, which lowers the physical gas densities in the absorbing systems and thereby lowers their neutral fractions and corresponding HI column densities. It is therefore more physically appropriate to think of $f(N_{\text{HI}})$ as evolving to the left rather than evolving downwards, though the quantitative effect is the same to within the accuracy of this simplified account.

Figure 3 shows the distribution of b -parameters for lines with $N_{\text{HI}} \geq 10^{13} \text{ cm}^{-2}$ from HKCSR (solid histogram) and from the AUTOVP analyses of the simulation at $z = 3$ and $z = 2$ (solid and dashed curves, respectively). The 10^{13} cm^{-2} cutoff eliminates lines whose identification and derived properties are sensitive to the value of S/N or to details of the fitting procedure, though the results do not change qualitatively if we lower this cutoff to $10^{12.5} \text{ cm}^{-2}$. The threshold method (dotted curve) yields much larger b -parameters than AUTOVP at $z = 3$ because many of its identified “lines” are extended absorption regions, which AUTOVP separates into narrower components. Table 1 lists the median, mean, and 1σ dispersion of the b -parameter histograms. The $z = 3$ simulation values for all three numbers are slightly larger than the HKCSR values, but the agreement is quite good given that the analysis procedures are not identical in all their details. The most significant difference in the distributions is the presence of many more narrow ($b < 20 \text{ km s}^{-1}$) lines in the simulation than in the data, a discrepancy that we discuss further below.

5. Discussion

Our most important result is that the CDM simulation reproduces the large number of weak lines found in HIRES spectra when it is analyzed by Voigt-profile decomposition. However, in the column density range $10^{13} \text{ cm}^{-2} \lesssim N_{\text{HI}} \lesssim 10^{14} \text{ cm}^{-2}$, the density of lines in the simulation at $z = 3$ is a factor of $1.5 - 2$ higher than found by HKCSR at $\bar{z} = 2.9$. Our simulation suffers from the inevitable limitation of finite numerical resolution, but the Ly α absorbers are usually large, smooth, low-overdensity structures, and we would in any case expect higher numerical resolution to increase the number of lines rather than decrease it. This excess of lines may therefore indicate a failure of the $\Omega = 1$, $\sigma_8 = 0.7$ CDM model, a tendency to produce too much small scale clumping at $z = 3$.

An alternative possibility, quite plausible at present, is that we have set the intensity of the UV background too low given our adopted value of Ω_b . As discussed in §2, we choose the background intensity in order to match the PRS determination of the mean Ly α optical depth, $\bar{\tau}_\alpha = 0.45$ at $z = 3$. The statistical uncertainties in this determination are fairly small, but there are systematic uncertainties in the required extrapolation of the quasar continuum into the Ly α forest region. We can match the HKCSR results if we increase the UV background intensity by a factor ~ 1.7 , thus lowering HI column densities by a similar factor and shifting the simulation result for $f(N_{\text{HI}})$ to the left. This increase lowers the mean optical depth to $\bar{\tau}_\alpha = 0.32$, which is well outside the 1σ range of PRS (figure 4) at the HKCSR mean redshift $\bar{z} = 2.9$ but is consistent with the PRS value at $z = 2.65$. It is somewhat *above* the value $\bar{\tau}_\alpha(z = 3) \sim 0.25$ found by Zuo & Lu (1993), who use a different data set and a different method of determining the quasar continuum. The uncertainty of our conclusions highlights the need for better observational determinations of $\bar{\tau}_\alpha(z)$, which plays a crucial role in normalizing the predictions of cosmological simulations. The mean optical depth depends on the distribution of b -parameters and on the small scale clustering of absorbers in addition to the column density distribution itself, so if $\bar{\tau}_\alpha$ is well known then the amplitude of $f(N_{\text{HI}})$ becomes an important independent test of the high-redshift structure predicted by a cosmological model.

A second important result of our comparison is that the $\Omega = 1, \sigma_8 = 0.7$ CDM model produces Ly α forest lines with typical b -parameters close to observed values. However, the simulation yields many more lines with $b < 20 \text{ km s}^{-1}$ than are found by HKCSR. A thermally broadened, single-temperature gas cloud produces a Ly α absorption line of width $b = 20(T/24,000 \text{ K})^{1/2} \text{ km s}^{-1}$. Low- b lines arise in the simulated spectra because much of the absorbing gas is at temperatures of 10^4 K or less, with its temperature determined by the balance between photoionization heating and adiabatic cooling (MWHK). We will need tests at higher resolution to check that these temperatures are not artificially low because the simulation misses entropy production in unresolved shocks, but because such shocks would have to be quite weak, we expect that the effect of missing them would be small.

Miralda-Escudé & Rees (1994) pointed out that the process of reionization can heat the IGM significantly if it occurs fast enough to prevent radiative cooling losses. Our equilibrium treatment of photoionization (KWH) implicitly suppresses this effect, and since circumstantial evidence suggests that HeII reionization may have occurred at $z \approx 3$ (Songaila & Cowie 1996), we could be underestimating the gas temperatures at this redshift. The dot-dash line of Figure 3 shows the b -parameter distribution obtained at $z = 3$ after adding 15,000 K to the temperatures of the SPH particles (a thermal energy equivalent to 4 Rydbergs per HeII photoelectron), reducing the UV background intensity by 2.46 to restore $\bar{\tau}_\alpha = 0.45$, then reextracting and reanalyzing spectra. Heating the gas eliminates the excess of low b -parameter lines, though it worsens the agreement with HKCSR at $b > 40 \text{ km s}^{-1}$. We will investigate other treatments of reionization heating in future work, though the effects will be difficult to pin down because they depend on uncertain details of reionization (Miralda-Escudé & Rees 1994). Model predictions for b -parameters should

be more robust at $z = 2$, since by this time much of the energy absorbed during HeII reionization at $z \gtrsim 3$ will have been lost to adiabatic cooling.

Sharper tests of cosmological models against the statistics of the Ly α forest can be obtained by expanding the redshift range of comparisons, by improving the determination of $\bar{\tau}_\alpha(z)$, and by applying AUTOVP to observational data, so that the analyses of simulated and observed spectra are identical in detail. We will also test models of the Ly α forest using alternative statistical measures to characterize spectra, for if the physical scenario that emerges from cosmological simulations is correct, then Voigt-profile decomposition provides at best a rough guide to the density and temperature profiles of the absorbing gas. At the high S/N and resolution of the HIRES data, an absorption feature with a single flux minimum often shows asymmetries or broad wings, requiring two or more Voigt-profile lines to provide an adequate fit (HKCSR). Pairs of lines with small velocity separations have strongly anti-correlated b -parameters, suggesting that many of these decompositions are not genuine physical blends (Rauch 1996). While results such as these can always be accommodated within a discrete “cloud” model by postulating just the right clustering properties, they more likely signify the breakdown of the Voigt-profile paradigm itself, revealing the origin of the Ly α forest in the diffuse, undulating gas distribution of the high-redshift universe.

We acknowledge the invaluable assistance of Chris Churchill and numerous stimulating discussions with Jordi Miralda-Escudé. We thank the authors of HKCSR for making their line list available. We also thank Renyue Cen for his timely refereeing and helpful comments. This work was supported in part by the PSC, NCSA, and SDSC supercomputing centers, by NASA theory grants NAGW-2422, NAGW-2523, NAG5-2882, and NAG5-3111, by NASA HPCC/ESS grant NAG 5-2213, and by the NSF under grants ATS90-18256, ASC 93-18185 and the Presidential Faculty Fellows Program.

REFERENCES

- Bahcall, J. & Cen, R. 1992 ApJ, 398, L81
- Blumenthal, G. R., Faber, S. M., Primack, J. R., & Rees, M. J. 1984, Nature, 311, 517
- Cen, R., Miralda-Escudé, J., Ostriker, J.P., & Rauch M. 1994, ApJ, 427, L9
- Churchill, C.W. 1996, UCSC Ph.D. Thesis
- Gunn, J.E. & Peterson, B.A. 1965, ApJ, 142, 1633
- Haardt, F. & Madau, P. 1996, ApJ, 461, 20 (HM)
- Hernquist, L.H. & Katz, N. 1989, ApJS, 70, 419
- Hernquist, L.H., Katz, N., Weinberg, D.H., & Miralda-Escudé, J. 1996, ApJ, 457, L51 (HKWM)
- Hu, E.M., Kim, T.S., Cowie, L.L., Songaila, A., & Rauch, M. 1995 AJ, 110, 1526 (HKCSR)
- Katz, N., Weinberg D.H., & Hernquist, L. 1996, ApJS, 105, 19 (KWH)
- Katz, N., Weinberg D.H., Hernquist, L., & Miralda-Escudé, J. 1996b, ApJ, 457, L57
- Lanzetta, K.M., Turnshek D.A., & Wolfe, A.M. 1987, ApJ, 322, 739
- Miralda-Escudé, J., Cen, R., Ostriker, J.P., & Rauch, M. 1996, ApJ, in press, astro-ph/9511013
- Miralda-Escudé, J. & Rees, M. 1994, MNRAS, 266, 343
- Miralda-Escudé, J., Weinberg, D. H., Hernquist, L., & Katz, N. 1997, in preparation (MWHK)
- Peebles, P. J. E. 1982, ApJ, 263, L1
- Petitjean, P., Webb, J.K., Rauch, M., Carswell, R.F., & Lanzetta, K. 1993, MNRAS, 262, 499 (PWRCL)
- Press, W.H. & Rybicki, G.B. 1993, ApJ, 418, 585
- Press, W.H., Rybicki, G.B., & Schneider, D.P. 1993, ApJ, 414, 64 (PRS)
- Rauch, M. 1996, in Cold Gas at High Redshift, eds. M. Bremer, H. Rottgering, C. Carilli, & P. van de Werf, (Dordrecht: Kluwer)
- Rugers, M. & Hogan, C. J. 1996, ApJ, 459, L1
- Songaila, A. & Cowie, L. L. 1996, AJ, 112, 335
- Tytler, D., Fan, X.M., & Burles, S. 1996, Nature, 381, 207

Vogt, S. S., et al. 1994, SPIE, 2198, 326

Weinberg, D. H., Hernquist, L., & Katz, N. 1996, ApJ, submitted, astro-ph/9604175

White, S. D. M., Efstathiou, G., & Frenk, C. S. 1993, MNRAS, 262, 1023

Zhang, Y., Anninos, P., & Norman, M.L. 1995, ApJ, 453, L57

Zuo, L. & Lu, L. 1993, ApJ, 418, 601

Fig. 1.— Example AUTOVP fit (smooth curves) for a continuum-normalized, $z = 3$, artificial spectrum (histograms). The bottom panel shows the first-guess fit, while the top panel shows the final fit after χ^2 -minimization. Long tick marks indicate lines with $N_{\text{HI}} \geq 10^{13}$, while short tick marks indicate lines with $N_{\text{HI}} < 10^{13}$.

Fig. 2.— Column-density distributions $f(N_{\text{HI}}) \equiv d^2N/dz\,dN_{\text{HI}}$, the number of lines per unit redshift per linear interval of HI column density. Solid and dashed lines show the AUTOVP results for 300 spectra along random lines of sight through the simulation at $z = 3$ and $z = 2$, respectively. The dotted line shows $f(N_{\text{HI}})$ obtained by the threshold method at $z = 3$. Filled and open circles show the observational results of PWRCL and HKCSR, respectively. Error bars denote 1σ Poisson counting errors.

Fig. 3.— Distributions of b -parameters, for lines with $N_{\text{HI}} > 10^{13}\text{cm}^{-2}$. Solid and dashed curves show the AUTOVP results from the simulation at $z = 3$ and $z = 2$, respectively. The dotted line shows the result of the threshold method at $z = 3$; blending of features leads to very large line widths. The dot-dash line shows the AUTOVP result at $z = 3$ after gas temperatures have been increased by 15,000 K. The solid histogram shows the HKCSR data. All distributions are computed as histograms with the same bins, but the simulation results are shown as curves to prevent visual confusion.

Table 1. Moments of b -parameter distributions.

Spectra	Algorithm	\bar{z}	b_{med}	b_{mean}	σ_b
HKCSR		2.9	31.4	35.8	16.3
CDM	AUTOVP	3	34.6	39.3	20.8
CDM	AUTOVP	2	30.6	34.7	16.5
CDM+15K	AUTOVP	3	40.9	43.8	19.5
CDM	Threshold	3	74.5	83.9	47.6
CDM	Threshold	2	43.4	49.9	26.3



